

**METHOD AND APPARATUS FOR SCHEDULING TRAFFIC TO MEET
QUALITY OF SERVICE REQUIREMENTS IN A COMMUNICATION NETWORK**

5 **Field of the Invention**

The present invention relates generally to communication networks, and more specifically to scheduling techniques for controlling the transfer of data through such networks.

10 **Background of the Invention**

15 A very important and challenging problem in the design of high-speed communication networks is that of providing Quality of Service (QoS) guarantees, usually specified in terms of loss probabilities or packet delays in the network. For example, the control of packet delays is often of crucial importance, particularly for real-time applications such as video delivery systems, wireless networks, multimedia networks, call centers, etc. A basic decision that has to be made in such contexts is that of connection admission control, i.e., one has to determine when a new user can be admitted to the system, while still fulfilling the QoS requirements of all users already in the system. Moreover, users already in the system have to be scheduled in the most efficient manner so as to maximize the number of users that can be admitted into the system. Unfortunately, conventional scheduling techniques are unable to provide adequate QoS guarantees.

20 A number of these conventional scheduling techniques are based on deterministic QoS guarantees, as opposed to probabilistic guarantees. For example, in L. Georgiadis, R. Guerin and A. Parekh, "Optimal multiplexing on a single link: delay and buffer requirements," IEEE Transactions on Information Theory, 43(5):1518-1535, 1997, and J. Liebeherr, D. Wrege, and D. Ferrari, "Exact admission control for networks with a bounded delay service," IEEE/ACM
25 Transactions on Networking, 4(6):885-901, 1996, it has been shown that an Earliest Deadline First (EDF) scheduling technique is optimal in the context of providing deterministic QoS guarantees for a single node of a communication network. Another known technique, Coordinated EDF (CEDF), described in M. Andrews, L. Zhang, "Minimizing End-to-End Delay in High-Speed Networks with a Simple Coordinated Schedule," IEEE INFOCOM'99, pp. 380-388, 1999, uses randomness to
30 spread out packet transmission deadlines.

It is generally believed that deterministic QoS requirements such as those associated with the above-noted conventional scheduling techniques can lead to an overly conservative admission policy, and a consequent decrease in system throughput. In contrast, probabilistic QoS requirements usually provide the very advantageous trade off of a small amount of QoS for a large capacity gain.

5 A need therefore exists for improved scheduling techniques that can meet probabilistic QoS requirements.

Summary of the Invention

The invention provides an improved scheduling policy, referred to herein as Largest

10 Weighted Delay First (LWDF) scheduling, which can be used to meet probabilistic QoS requirements in a communication network. In accordance with the invention, a link scheduler associated with a node of a communication network selects packets for transmission over a communication link of the network. The link scheduler computes a delay measure W_i , $i = 1, 2, \dots$

15 $\dots N$, for each of N packets, where each of the N packets is associated with a corresponding one of N data flows and located in a head position in a corresponding one of N data flow queues in the network node. The delay measure for a given one of the packets may be computed, e.g., as a difference between a current time and an arrival time of the given packet in its corresponding queue in the network node. The link scheduler weights the computed delay measures using a set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$, and then selects for transmission the particular packet having the largest

20 weighted delay associated therewith.

In an illustrative embodiment of the invention configured to meet a QoS requirement specified in terms of a transmission deadline time T_i and an allowed deadline violation probability δ_i , e.g., a requirement specified by $P(W_i > T_i) \leq \delta_i$, the weights α_i in the set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$ may be given by $\alpha_i = -T_i / \log \delta_i$.

25 Advantageously, the invention can also be used in a manner similar to that described above to meet a variety of other types of probabilistic QoS requirements, including, e.g., requirements based on packet loss probabilities. For example, the QoS guarantee may be defined for a delay measure in the form of queue length Q_i , $i = 1, 2, \dots N$, and an allowed queue length violation probability δ_i . In such an embodiment, the QoS requirement is specified by $P(Q_i > H_i) \leq \delta_i$, and the

weights α_i in the set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$ may be given by $\alpha_i = -H_i / \log \delta_i$, where H_i represents an upper bound on the length of the queue.

The LWDF scheduling in the above-noted illustrative embodiments of the present invention is optimal in the sense that it maximizes the asymptotic rate of decay of the tails of a stationary maximal-weighted delay or queue length within a general class of scheduling policies. Thus, whenever it is feasible to satisfy specified QoS constraints, one would expect that LWDF would do so. However, in cases when it is not feasible, the LWDF scheduling policy of the illustrative embodiment will most likely violate the QoS constraints of most users. This has significant implications for flow admission control. Most importantly, LWDF in the illustrative embodiment allows one to detect in real time the infeasibility of satisfying QoS requirements. In addition, even when the QoS constraints are infeasible, LWDF in the illustrative embodiment has the property of fairness, in that it tries to equalize for all users the ratio of the logarithm of actual violation probability to the logarithm of the desired probability.

The LWDF scheduling of the invention provides a number of other advantages. For example, the form of the algorithm is invariant to changes in stochastic input flow models. In addition, it is suitable for use in a wide variety of communication network applications, including wireless networks, high-speed multimedia networks, call centers, etc.

The scheduling techniques of the present invention may be implemented at least in part in the form of one or more software programs running on a computer or other type of programmable digital processor.

Brief Description of the Drawings

FIG. 1 shows an illustrative embodiment of the invention as implemented in a portion of a communication network

FIG. 2 is a flow diagram of the operation of a link scheduler in the illustrative embodiment of FIG. 1.

FIG. 3 shows a block diagram of an exemplary communication network in which the invention may be implemented.

Detailed Description of the Invention

The present invention will be illustrated below in conjunction with exemplary scheduling techniques implemented in a link scheduler associated with a particular node of a communication network. It should be understood, however, that the invention is not limited to use with any particular type of communication network device or network configuration. The disclosed techniques are suitable for use with a wide variety of other systems and in numerous alternative applications. The term “packet” as used herein is intended to include any grouping or other arrangement of data suitable for transmission over a communication network link. A given “packet” in accordance with this definition may therefore include only a single bit. The term “delay measure” as used herein is intended to include not only time-based measures, but also other types of measures indicative of delay, such as, e.g., queue length measures.

An illustrative embodiment of the invention optimally schedules classes of heterogeneous users in a communication network in order to meet stipulated QoS requirements for each class, as specified in terms of a deadline T_i and an allowed violation probability δ_i . More specifically, assuming there are N classes of users and W_i is a delay associated with a packet of the i th class, then the illustrative embodiment provides an optimal scheduling policy that meets the given QoS constraints

$$P(W_i > T_i) \leq \delta_i \quad \text{for } i = 1, \dots, N. \quad (1)$$

This can alternatively be expressed as an optimal scheduling policy which maximizes the quantity

$$\min_{i=1, \dots, N} \frac{\log P(W_i > T_i)}{\log \delta_i}. \quad (2)$$

As will be described in greater detail below, feasibility of the QoS constraints can then be determined by analyzing the performance of this optimal policy.

FIG. 1 shows an illustrative embodiment of the invention as implemented in a node 10 of a communication network. The node 10, also referred to as node A, receives a set of input data flows 11 from, e.g., other nodes or information sources of the network. The node 10 communicates with another node 12, also referred to as node B, via an A→B communication link 14. The nodes 10, 12 and communication link 14 represent portions of the communication network, e.g., nodes 10, 12 may represent computers or other types of digital data processing devices, and link 14 may represent a packet-based transmission medium over which the nodes communicate.

The node 10 as shown in FIG. 1 has N input data flows, each of which is associated with a corresponding class of users and is applied to a corresponding packet queue 15-1, 15-2, . . . 15- N within the node 10. It is assumed without limitation that in the illustrative embodiment each data flow has stationary increments and satisfies a sample path large deviation principle (LDP). LDP techniques are described in greater detail in, e.g., A. Dembo and O. Zeitouni, "Large Deviations Techniques and Applications," 2nd Edition, Springer, 1998, and A. Shwartz and A. Weiss, "Large Deviations for Performance Analysis: Queues, Communication and Computing," Chapman and Hall, New York, 1995. It should be emphasized that the invention does not require these particular assumptions, i.e., embodiments of the invention may be constructed for which the above assumptions and other assumptions mentioned herein do not hold.

The node 10 further includes a link scheduler 16 which retrieves packets from the queues 15-1, 15-2, . . . 15- N for transmission over the link 14. In accordance with the invention, the link scheduler 16 implements a Largest Weighted Delay First (LWDF) scheduling policy based on a set of positive weights $\alpha = \alpha_1, \alpha_2, \dots, \alpha_N$. Generally speaking, the LWDF scheduling policy associated with the vector α is a non-preemptive policy that always chooses for transmission the user class for which, e.g., the quantity W_i / α_i is maximum, where W_i is the delay experienced by the head-of-the-queue packet of user class i , i.e., the delay experienced by the packet at the head of the queue 15- i in node 10. When all user classes i have the same QoS requirements, the LWDF policy reduces to a well-known conventional first-in first-out (FIFO) policy. Advantageously, the form of the scheduling algorithm in this illustrative embodiment is invariant to changes in stochastic input flow models.

Under the above-noted assumptions that the input flows have stationary increments and satisfy a sample path LDP, and with the further assumption that the system is stable, it can be shown that the LWDF policy of the present invention is asymptotically optimal in the sense that, within a large class of so-called work-conserving policies, it maximizes the quantity

$$\min_{i=1,\dots,N} \left[\alpha_i \lim_{n \rightarrow \infty} \frac{-1}{n} \log P(W_i > n) \right]. \quad (3)$$

The above criterion can be related back to the original QoS requirement (2) by choosing the LWDF weight vector α such that $\alpha_i = -T_i / \log \delta_i$ for $i = 1, \dots, N$. Then the above optimality result shows that the LWDF policy of the invention corresponding to this vector α maximizes the quantity (2) above within a large class of scheduling policies, in the limit as deadlines T_i get large and acceptable violation probabilities δ_i get small. In other words, if it is at all feasible to meet the asymptotic QoS constraints (3), then the LWDF policy of the invention will do so. To prove this optimality result, it can be shown that

$$\min_{i=1,\dots,N} \left[\alpha_i \lim_{n \rightarrow \infty} \frac{-1}{n} \log P(W_i > n) \right] = \lim_{n \rightarrow \infty} \frac{-1}{n} \log P \left(\max_{i=1,\dots,N} \frac{W_i}{\alpha_i} > n \right). \quad (4)$$

Thus, the problem of maximizing (3) is reduced to that of maximizing the quantity on the right hand side of (4) above, which can be analyzed using large deviation techniques such as those described in the above-cited LDP references. If the maximal stationary weighted delay \hat{r} is defined to be

$$\hat{r} = \max_{i=1,\dots,N} W_i / \alpha_i,$$

then an LDP can be derived for the sequence of scaled maximal weighted delays. In particular, the quantity (3) under the LWDF scheduling policy can be characterized by a finite-dimensional variational problem. Moreover, an LDP can be obtained for each sequence of scaled stationary delays W_i , under the LWDF scheduling policy, and the associated rate functions can be characterized in terms of finite-dimensional variational problems. It should be noted, however, that the invention can be implemented using other types of weights, i.e., weights other than the particular example of α_i given above.

FIG. 2 is a flow diagram illustrating the manner in which the above-described LWDF scheduling policy may be implemented in the link scheduler 16 of FIG. 1. In step 20, a new packet arrives in node 10. The packet is time stamped based on its arrival time in the node, as shown in step 22. A determination is then made in step 24 as to whether the link scheduler 16 is in a "waiting state," i.e., a state in which the scheduler is waiting for a new packet to arrive. If the scheduler is not in this state, then no action is taken, as indicated by step 26. If the scheduler is in the waiting state, a determination is made in step 28 as to whether there are packets waiting for transmission. If there are no packets waiting for transmission, the scheduler in step 30 waits for a new packet to arrive and enters the waiting state.

If there are packets waiting for transmission, the scheduler in step 32 computes the delay W_i for each data flow i , $i = 1, 2, \dots, N$. This delay may be computed for a given data flow, e.g., by taking the current time and subtracting the arrival time of the packet at the head of the corresponding queue 15- i . If a given queue includes no packets, it could be automatically dropped from consideration, or its W_i value could be set to zero. The scheduler in step 34 then determines j , where $j = 1, 2, \dots, N$, such that the value W_j/α_j is maximal among all W_i/α_i for $i = 1, 2, \dots, N$, where α_i is given by $\alpha_i = -T_i / \log \delta_i$ as previously described. The scheduler 16 then transmits the packet at the head of the j th queue, i.e., the head-of-the-line packet at queue 15- j . The process of steps 32, 34 and 36 is repeated as long as there are packets waiting for transmission. If two or more j values in step 34 yield the same result, a particular one of the j values may be selected arbitrarily, e.g., the largest of the multiple j values producing the maximal W_j/α_j may be selected.

FIG. 3 shows a portion of an exemplary communication system 100 in which the above-described LWDF scheduling policy may be implemented. The system 100 includes a plurality of

nodes 102 which communicate over a network 104. The network 104 may itself be comprised of multiple interconnected nodes, and may represent, e.g., the Internet, a wireless network, a local area network, a wide area network, an intranet, an extranet, a telephone, cable or satellite network, as well as combinations or portions of these and other networks. A given one of the nodes 102, denoted
 5 node A, corresponds generally to node A of the FIG. 1 configuration. Node A in this example includes a processor 110 and a memory 112.

The processor 110 may represent, e.g., a microprocessor, a computer, a central processing unit (CPU), an application-specific integrated circuit (ASIC), as well as portions or combinations of these and other suitable processing devices. The memory 112 may represent, e.g., an electronic
 10 memory, a magnetic or optical disk-based memory, as well as portions or combinations of these and other memory devices. For example, the memory 112 may represent a memory associated on a common integrated circuit with the processor 110, or a separate electronic memory integrated with the processor 110 into a computer or other device, and may be used to implement the queues 15-*i* of FIG. 1. The link scheduler 16 of FIG. 1 may be implemented in whole or in part in the form of
 15 one or more software programs stored in the memory 112 and executed by the processor 110.

The LWDF approach in the above-described illustrative embodiment of the invention is optimal in the sense that it maximizes the asymptotic rate of decay of the tails of the stationary maximal weighted delay within a general class of scheduling policies. For large delays and small
 20 allowed violation probabilities, the LWDF policy with weights $\alpha_i = -T_i / \log \delta_i$ is a nearly optimal policy to use in order to satisfy the QoS constraints given by (1). Thus, whenever it is feasible to satisfy these constraints, one would expect that LWDF would do so. However, in cases when it is not feasible, the LWDF policy will most likely violate the QoS constraints of most users. This has significant implications for flow admission control. Most importantly, LWDF allows one to detect
 in real time the infeasibility of satisfying QoS requirements. In addition, even when the QoS
 25 constraints are infeasible, LWDF has the property of fairness, in that it tries to equalize for all users the ratio of the logarithm of actual violation probability to the logarithm of the desired probability. However, this notion of fairness may or may not be desirable in a particular application. For example, in certain applications, it may be preferable to maintain feasibility of as many users as

possible, while penalizing the rest. In such applications, one may want to use the LWDF of the invention for only a subset of flows, while giving lower priority to other flows.

The LWDF scheduling of the invention provides a number of other advantages. For example, it is invariant to changes in stochastic input flow models. In addition, it is suitable for use in a wide variety of communication network applications, including wireless networks, high-speed multimedia networks, call centers, etc.

It should be noted that the above-described results concern stationary delays, rather than stationary waiting times of individual users. Nevertheless, one would expect that the asymptotics of the tails of stationary distributions of both these processes typically coincide. It is therefore believed that similar results would be obtained for stationary waiting times.

The above-described embodiments of the invention are intended to be illustrative only. For example, the invention can be implemented in a wide variety of different data processing devices, using software, hardware or combinations of software and hardware. In addition, other types of delay measures and weights can be used to implement the LWDF scheduling policy. Moreover, the invention may be used in conjunction with other types of QoS guarantees, including, e.g., guarantees based on packet loss probability. For example, the QoS guarantee may be defined for a delay measure in the form of queue length Q_i , $i = 1, 2, \dots, N$, and an allowed queue length violation probability δ_i . In such an embodiment, the QoS requirement is specified by $P(Q_i > H_i) \leq \delta_i$, and the weights α_i in the set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$ may be given by $\alpha_i = -H_i / \log \delta_i$, where H_i represents an upper bound on the length of the queue. These and numerous other alternative embodiments may be devised by those skilled in the art without departing from the scope of the following claims.